



US006477147B1

(12) **United States Patent**
Böcking et al.

(10) **Patent No.: US 6,477,147 B1**
 (45) **Date of Patent: Nov. 5, 2002**

(54) **METHOD AND DEVICE FOR
 TRANSMITTING A DATA PACKET USING
 ETHERNET FROM A FIRST DEVICE TO AT
 LEAST ONE OTHER DEVICE**

(75) Inventors: **Stefan Böcking, München; Bodo
 Braun, Höhenkirchen/Siegersbrunn;
 Hans-Peter Huth, München, all of
 (DE)**

(73) Assignee: **Siemens Aktiengesellschaft, Munich
 (DE)**

(*) Notice: Subject to any disclaimer, the term of this
 patent is extended or adjusted under 35
 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/142,116**

(22) PCT Filed: **Mar. 4, 1997**

(86) PCT No.: **PCT/DE97/00404**

§ 371 (c)(1),
 (2), (4) Date: **Sep. 4, 1998**

(87) PCT Pub. No.: **WO97/33408**

PCT Pub. Date: **Sep. 12, 1997**

(30) **Foreign Application Priority Data**

Mar. 8, 1996 (DE) 196 09 080
 Mar. 8, 1996 (EP) 96130699

(51) Int. Cl.⁷ **G01R 31/08; H04O 7/24**

(52) U.S. Cl. **370/236; 370/338**

(58) Field of Search 370/236, 400,
 370/428, 429, 414, 413, 412, 363, 368,
 371, 373, 374, 378, 381

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,815,110 A 3/1989 Benson et al.

5,319,641 A 6/1994 Fridrich et al.
 5,353,287 A 10/1994 Kuddes et al.
 5,600,641 A * 2/1997 Duault et al. 370/400
 5,724,358 A * 3/1998 Headrick et al. 370/418
 5,790,522 A * 8/1998 Fichou et al. 370/236
 5,936,958 A * 8/1999 Soumiya et al. 370/395
 5,996,019 A * 11/1999 Hauser et al. 709/235
 6,219,728 B1 * 4/2001 Yin 710/52

FOREIGN PATENT DOCUMENTS

EP 0 333 225 A2 9/1989
 WO WO 87/01545 4/1987

OTHER PUBLICATIONS

UNIX Internetworking, Uday O. Pabrai, Artech House,
 Boston, London.

Rate-Controlled Static-Priority Queueing, Hui Zhang et al,
 Computer Science Division, 1993 IEEE, University of Cali-
 fornia at Berkley.

* cited by examiner

Primary Examiner—Wellington Chin

Assistant Examiner—Brenda Pham

(74) *Attorney, Agent, or Firm*—Staas & Halsey LLP

(57) **ABSTRACT**

A method and arrangement for transmitting a data packet in
 the Ethernet from a first arrangement to at least a second
 arrangement.

When a data packet is transmitted in the Ethernet, the data
 packet is assigned a priority class (201). In addition, the data
 packet is labeled (202) in accordance with the priority class
 and transmitted (203). As a result, it becomes possible also
 to transmit in the Ethernet data packets of a communications
 connection with real-time requirements, it being possible to
 ensure specific time requirements in the statistical sense.

10 Claims, 5 Drawing Sheets

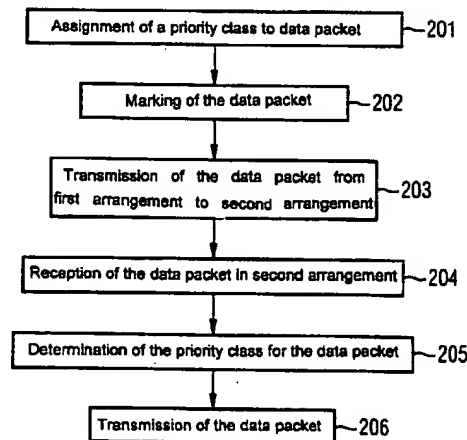


FIG 1


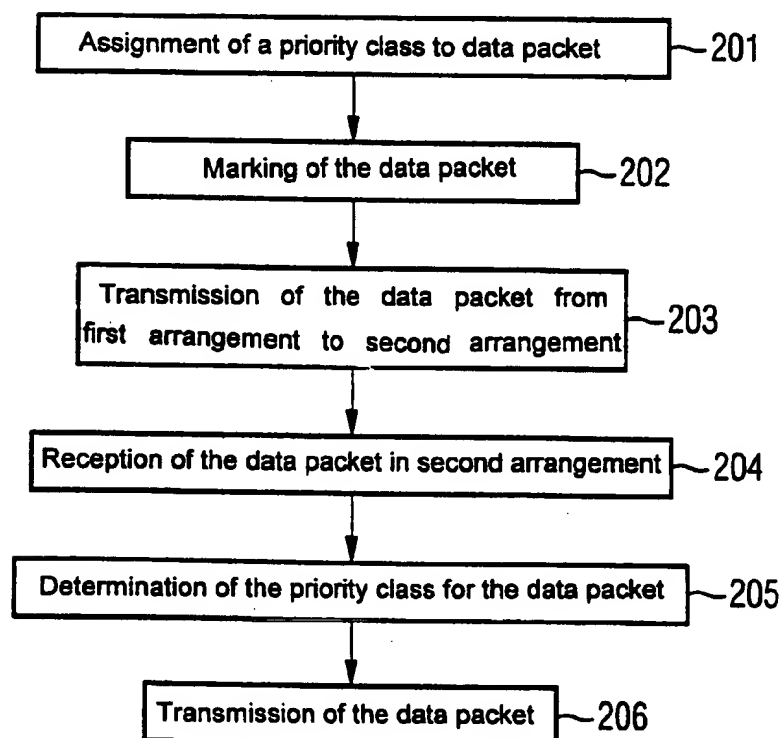


FIG 2








FIG 4






FIG 6

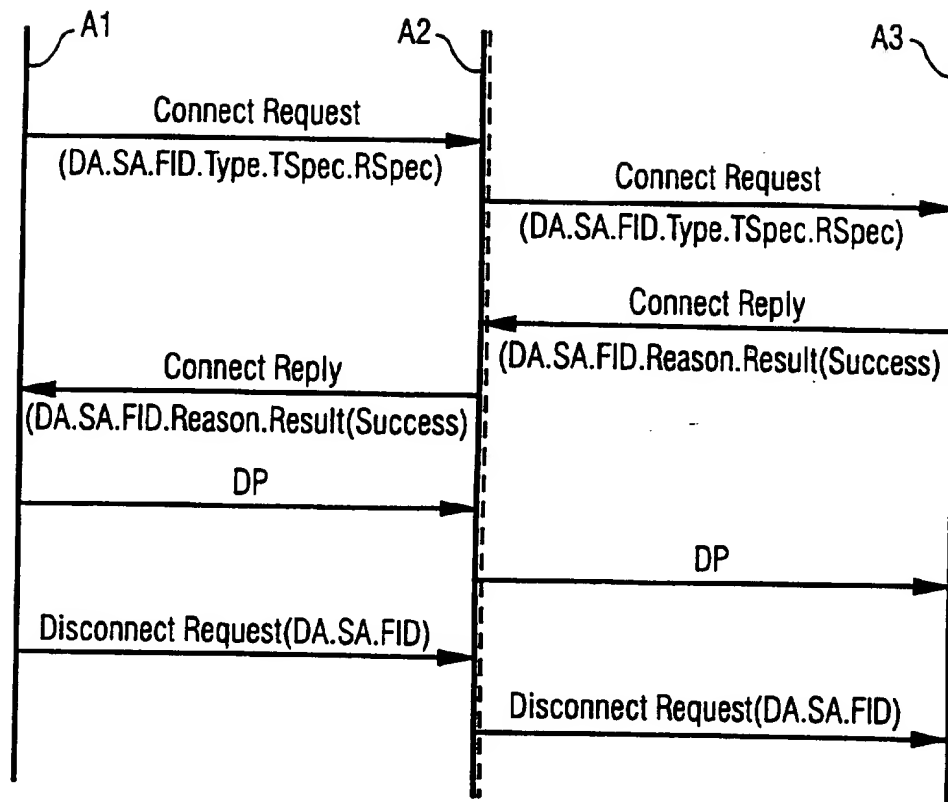



FIG 7



METHOD AND DEVICE FOR TRANSMITTING A DATA PACKET USING ETHERNET FROM A FIRST DEVICE TO AT LEAST ONE OTHER DEVICE

BACKGROUND OF THE INVENTION

Local computer networks which are configured according to the IEEE 802.3 Standard, which is referred to below as Ethernet Standard, constitute a technology with which terminals are connected via a commonly used, serial bus. The access to this bus is regulated by means of the so-called Carrier Sense Multiple Access/Collision Detection (CSMA/CD). The Ethernet protocol is a so-called fair protocol, i.e. when there is competing access, all the terminals connected to a bus have, on a statistical average, the same chance of transmitting via the bus their data which is to be transmitted, regardless of the type of data to be transmitted.

However, in many cases it is necessary to prefer certain data flows over others. Examples of this are data streams with real-time requirements, for example a sound data stream or a video data stream, but also alarm messages for controlling machines. Such data streams require a certain minimum quality during their transmission. Guarantees for such minimum quality levels are, however, currently not ensured with the Ethernet protocol, since with the CSMA/CD protocol it is not possible to make a distinction between data streams and therefore different treatment of data streams is also impossible.

The Ethernet bus is a passive transmission medium and the switching work is distributed between the terminals which also operate according to the Ethernet protocol. The so-called Switched Ethernet is a technology in which local networks according to the IEEE 802.3 Standard are coupled by means of packet switching instead of the otherwise customary Ethernet bus. Packet switching is implemented by means of the so-called Ethernet node. An Ethernet node has a plurality of bidirectional accesses, the so-called ports. The Ethernet node must pass on incoming data packets through at least one output port which is coupled directly or indirectly to the destination or destinations of the data packet. If different input packets are received and have to be output via the same output port, the packets are buffered. The buffer used for this can overflow in cases of high load and newly arriving data packets are lost in this case.

Various packet formats of the Ethernet protocol are known from document by U. O. Pabrai, UNIX Internetworking, Artech House, Boston, London, page 23, 1993.

SUMMARY OF THE INVENTION

The invention is based on the problem of specifying a method which is compatible with the Ethernet Standard and with which a certain minimum quality of the transmission for data streams is ensured with respect to real-time requirements for the transmission.

In the method, a priority class is assigned to each data packet to be transmitted by a first arrangement, and the data packet is marked according to the priority class. Then, the data packet is transmitted to a second arrangement taking into account the priority class.

A considerable advantage of the method is especially the fact that it is possible to prioritize the data packets, for example depending on the type of the data stream to be transmitted. In this way, it is possible, on a statistical

average, to ensure the required quality for data packets which require a higher priority because of quality requirements relating to real-time requirements during transmission.

The expression "on a statistical average" is to be understood in this context to mean that it is possible to ensure the quality requirements with a certain degree of probability according to the priority class used. The cause of this is that the data packets are transmitted according to their priority, for example data packets with a relatively high priority are preferred over data packets with a relatively low priority.

In order to be able to implement a multi-stage communications connection, i.e. a communications connection via several Ethernet switching nodes, it is advantageous in one development of the invention that the data packet is received from a second arrangement, the priority class assigned to the data packet is determined and in the case of which, in turn, the data packet is transmitted on taking into account the priority class.

By means of this procedure, the method is simplified since an entirely new priority class for the data packet does not, in turn, have to be formed in each switching node and assigned to the data packet, rather merely the priority class which has previously been assigned to the data packet is determined, and the priority class continues to be used in the rest of the method.

In order to assign the priority class to the data packet it is advantageous to analyze information contained in the data packet fed to the so-called Ethernet layer and relating to the type of the data packet and thus to the type of the communications connection or else to the type of the data stream, and to take into account the analyzed information in the assignment of the priority class. In this way, it becomes possible automatically to analyze and to ensure the type of the data packet, and thus the quality requirements of the data stream.

In addition, in one development of the invention it is advantageous to carry out an access check for the data packet, as a result of which it is possible to prevent the arrangement with which the method is carried out from being overloaded as a result of an excessive number of data packets to be transmitted.

In addition, in one development of the invention it is advantageous to divide up a buffer of the arrangement into a plurality of parts and to assign each part of the buffer to at least one priority class in each case. This means that in each case only data packets of the corresponding priority class can be stored in the corresponding part of the buffer.

The development described above is improved even more by virtue of the fact that the data packets are output from the corresponding parts of the buffer and transmitted in a predefinable sequence. The sequence can be predefined according to any desired scheduling method.

In addition it is advantageous that, in the event of the buffer memory, or part of the buffer memory in an arrangement which is configured as an Ethernet switching element, overflowing, the data packet causing the overflow is rejected only if it is not possible to abort the reception procedure by means of an artificial collision for this data packet.

In addition, in one development it is advantageous to generate the artificial collisions of the data packets only for data packets for whose source arrangement the generation of an artificial collision is known to be permissible, i.e. can be processed by the transmitting arrangement, for example. This development makes a selection for the arrangements which can process an artificial collision of a transmitted data packet.

In addition, in one development it is advantageous to reserve, during the entire communications connection, resources which are negotiated in a connection setup phase and are required during the communications connection, for data packets which are assigned to a prescribable communications connection. In this way, considerably more reliable guarantees for the quality levels which are required of the data packets to be transmitted are fulfilled.

The arrangement for carrying out the method has a data-packet classification unit, a priority-marking unit and a buffer memory. Using the data-packet classification unit, the data packet is assigned the priority class, and using the priority-marking unit the data packet is unambiguously labeled in accordance with the priority class. The data packets to be transmitted are buffered in the buffer.

In addition, in one development it is advantageous to provide a resource-administration unit in the arrangements and the switching node, which unit has at least one of the following functions:

- administration of resources,
- access control for the data packet and/or
- signaling between the coupled arrangements.

By virtue of the resource-administration unit, a simplified organization of the execution of the method is possible, enabling the method to be improved.

In addition, it is advantageous to assign the resource-administration unit of the switching node its own so-called media access control address (MAC address), as a result of which it becomes possible to address the resource-administration unit separately and thus also to implement it independently of the arrangement which carries out the method.

BRIEF DESCRIPTION OF THE DRAWINGS

The features of the present invention which are believed to be novel, are set forth with particularity in the appended claims. The invention, together with further objects and advantages, may best be understood by reference to the following description taken in conjunction with the accompanying drawings, in the several Figures of which like reference numerals identify like elements, and in which:

FIG. 1 shows a diagram of arrangements of a communications connection with two terminals and a switching unit, which exchange data packets according to the Ethernet protocol;

FIG. 2 shows a flowchart in which the method is illustrated in its method steps with some developments;

FIGS. 3a to 3f show a diagram in which various protocol formats of the Ethernet Standard, and various associated possible ways of marking a data packet are illustrated;

FIG. 4 shows a diagram of the arrangement, with a number of further developments and detailed illustrations with which the interaction with further communications layers is described;

FIGS. 5a to 5f show various protocol formats and associated possible ways of implementing a back-pressure method (described below) in compliance with the Ethernet Standard;

FIG. 6 shows a message flowchart in which, by way of example, a signaling operation and the exchange of data of the arrangements illustrated in FIG. 1 is described;

FIG. 7 shows a diagram in which a chronological division of the bandwidth of the communications connection into a plurality of time intervals for various priority classes is illustrated.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

FIG. 1 illustrates arrangements of a communications connection (illustrated by way of example), in which data packets are transmitted according to the Ethernet Standard.

A first arrangement A1, a second arrangement A2 and further arrangements Ai are illustrated.

In this exemplary case, the first arrangement A1 is embodied as an Ethernet terminal. The second arrangement A2 is embodied as a switching element, as a so-called Ethernet switching node (Ethernet switch). A further arrangement Ai is in turn configured in this exemplary case as an Ethernet terminal. Each further arrangement Ai is unambiguously marked with a first index i, the first index i being any desired natural number.

Between the arrangements A1, A2, Ai there is a point-to-point connection, through a first coupling K1 between the first arrangement A1 and the second arrangement A2, and a second coupling K2 between the second arrangement A2 and the further arrangement Ai.

However, it is also possible for arrangements which are configured as Ethernet terminals to be coupled directly to one another in the so-called shared Ethernet, and not separated by means of at least one arrangement configured as a switching element, as in the so-called switched Ethernet.

Since the Ethernet Standard is oriented according to the so-called layer model of the International Standardization Organization (ISO), the arrangements in the main illustration are described in the form of communication layers. The couplings K1, K2 form the physical medium, together with the so-called physical layer PHY, which may also be referred to as a bit transmission layer. The Ethernet layer, for example the layer of Media Access Control (MAC layer) is coupled to the bit transmission layer. The standardized Ethernet protocol is carried out in the Ethernet layer. The expansion of the Ethernet Protocol according to the invention can be considered logically as an intermediate layer between the MAC layer and further, higher communications layers. For this reason, in FIG. 1, the method is illustrated as an independent layer RMAC (Real Time Media Access Control).

The independent layer RMAC designated below as intermediate layer RMAC is, on the one hand, coupled to the MAC layer and, on the other hand, coupled to further communications layers. Further communications layers are to be understood as any desired, known transportation protocols, for example TCP/IP (Transport Control Protocol/Internet Protocol) or UDP/IP (User Datagram Protocol/Internet Protocol) or IPX, etc. All the communications layers which are arranged logically above the intermediate layer RMAC are referred to in their totality in FIG. 1 as higher layers HL for sake of simplicity. The general structure of the arrangements relating to the higher layers HL and of the bit transmission layer PHY and the MAC layer can be as desired within the scope of the respective protocol being used. Since the second arrangement in this exemplary case is configured as an Ethernet switch, the second arrangement does not have any higher layers HL in this case. This is not necessary according to the Ethernet Standard.

Even if only three arrangements are illustrated in FIG. 1, a communications connection in which the method is carried out can extend over any desired number of arrangements since the Ethernet protocol relates only to a connection of arrangements which are directly coupled to one another.

The standardized Ethernet protocol has, in particular, the disadvantage that, as a result of the free access, in accor-

5

dance with the CSMA/CD protocol, of all the equipment connected to the Ethernet bus, free access is possible for all the data packets DP to be transmitted, irrespective of the type of data to be transmitted.

The result of this is that for data streams which make predefinable real-time requests relating to the transmission of the data packets DP of the corresponding communications connection, [lacuna] cannot be ensured.

The various real-time requests are referred to below as quality features. Quality features are to be understood, within the scope of this document, as, for example, the following criteria:

- delay times of the data packets DP,
- delay time fluctuations in the data packets DP,
- losses of the data packets DP given overloading of the communications connection or of the arrangements,
- useful data rates, etc.

In order to be able to ensure specific quality features which are completely application-specific and be freely predefined in an application-specific fashion, the method has the following method steps which are illustrated in FIG. 2 together with some developments of the method.

The data packet DP to be transmitted is fed to the intermediate level RMAC by a higher level HL. In order to transmit the data packet DP in the Ethernet, the following method steps are provided.

In a first step 201, the data packet DP is assigned a priority class PKj from a quantity of any desired number of priority classes PKj. Each priority class PKj is unambiguously marked with the second index j, the second index j being any desired natural number.

In a further step 202, the data packet DP is unambiguously marked in accordance with the priority class PKj.

In a last step 203, the data packet DP is transmitted by the first arrangement A1, in which the method is being carried out, to the second arrangement A2.

In one development of the method, further method steps are provided which are also illustrated in FIG. 2.

In a further method step 204, the data packets DP is received by the second arrangement A2.

Then, in the second arrangement A2, the priority class PKj which has been assigned to the data packet DP in the first arrangement A1 is determined 205.

Taking into account the priority class PKj which has been determined, the data packet DP is transmitted on 206, in this exemplary case which is illustrated in FIG. 1, to the further arrangement Ai.

The concrete embodiment of the individual method steps in this exemplary embodiment will be explained in more detail below.

Assignment of the Priority Class PKj to the Data Packet DP 201

Generally, the assignment of the priority class PKj to the data packet DP signifies a mapping of the possibly large number of required quality features onto the priority class PKj, by means of which the corresponding combination of necessary quality features is ensured.

Even if the number and properties of the quality features are random, it has proven advantageous to take into account the following quality of service parameters (QoS parameters) in order to ensure the quality features in the method. It has proven sufficient for the following four priority classes Pkj (j=0, 1, 2, 3) to be taken into account within the scope of this method.

A first priority class PK0 is provided for a connectionless communications connection in which all that is assured is

6

that the respective data packet DP is transmitted, with an unspecified bit rate, as well as possible depending on the utilization rate of the communications network without anything being ensured and without a connection setup. Thus, the first priority class PK0 corresponds to the lowest priority which can be assigned to the data packet DP.

A second priority class PK1 is provided for a connection-orientated communications connection in which a controlled delay, i.e. a maximum delay of the transmission of the data packet DP, is ensured statistically. In the so-called controlled delay service, the following features, for example, are implemented:

- a bandwidth which is required for the connection and which is negotiated in a further described connection setup is ensured;
- the average delays during the transmission of the data packet DP when there is a high degree of loading of the communications network are in no way worse than the delays of the data packet DP to which the first priority class PK0 has been assigned;
- the maximum delay in the transmission of the data packet DP with a high degree of loading of the communications network are considerably lower than that of data packets DP of the first priority class PK0;
- the loss rate of the data packets DP owing to overflowing of the buffer memory PS of the arrangements A1, A2, Ai are not considerable as long as the properties of the data stream, that is to say of the communication connection negotiated in the "traffic contract" negotiated in the connection setup phase, for the respective communications connection are maintained.

With the second priority class PK1, a service is thus implemented with which a bandwidth demand with few real-time requirements, or with none at all, is guaranteed. The second priority class PK1 is suitable, for example, for burst-type traffic with a demand for a communications connection for a specific bandwidth.

A third priority class PK2 relates to a connection-orientated communications connection. The following services, for example are provided for data packets DP to which the third priority class PK2 has been assigned:

- most data packets DP are actually transmitted completely;
- in most cases, the delay of the data packets DP during the transmission will not exceed a predefinable maximum delay time for the transmission of the data packets DP.

In this context, the expression "most data packets" is to be understood such that it is a predefinable number which is to be stated, for example, in the connection setup phase. It is, for example, sufficient in many cases that one of 1000 transmitted data packets DP within a second and at maximum one of 10,000 data packets DP in even longer time intervals would exceed the limit given by the value for "most data packets". Thus, the third priority class PK2 will ensure a service in which a controlled delay plus a maximum delay limit is ensured, thus implementing a "quasi-real-time" service.

A fourth priority class PK3 is also provided for a connection-orientated communications connection. In the fourth priority class PK3 which corresponds to the highest priority which can be assigned to the data packet DP, the following services, for example are ensured:

- the bandwidth, for example, which is required for the communications connection and which is negotiated during the connection setup is made available;
- predefinable maximum delay times for the data packet DP are guaranteed during the transmission;

no losses of data packets DP occur owing to overflowing of the buffer PS.

Thus, the communications-connection parameters in the fourth priority class PK3 which are negotiated during the connection setup are ensured with a very much higher degree of statistical reliability than in the case of the other priority classes PK0, PK1, PK2 as long as the entire communications network does not collapse, for example owing to a fault in the communications network.

Marking of the Data Packet DP According to the Assigned Priority Class PKi 202.

The method of marking the data packet DP so that the receiver of the data packet DP can in each case determine which priority class PKi has been assigned to the data packet DP can be carried out in various ways.

A customary Ethernet data packet DP has, for example, the following fields in the data packet DP which are mentioned in the Ethernet Standard (cf. FIG. 3a):

- a destination address field DA in which the address of the receiver of the data packet DP is stated;
- a transmitter address field SA in which the address of the transmitter of the data packet DP is stated;
- a field type which is interpreted by the receiver of the data packet DP as an integer and is usually of the length of two octets. The field type is interpreted differently depending on the number in the field. If the number is smaller than 1500, the field type is interpreted, for example, as a long field, and the rest of the data packet DP corresponds to the customary so-called logic-link control format (LLC format). If the number in the field type is, however, not smaller than 1500, the number is interpreted as a type statement. The type statement contains a code for the network protocol which has generated the following information contained in the data packet DP.

In the so-called LLC format, the network protocol is defined by the so-called service access point (SAP), the LLC service access point (LLC SAP). An example of such a packet format is illustrated in FIG. 3c. FIG. 3e illustrates a variant of the LLC format, the so-called LLC subnet attachment point format (SNAP) in which the service access point (SAP) is set permanently to the value hexadecimal 0A. The network protocol is coded in a separate field type in this variant (cf. FIG. 3e).

In addition, all the protocol elements have an information field Info containing the actual information in the data packet DP which has been supplied by higher layers HL and is to be transmitted.

Also, a fault detection field FCS with a check sum code in the protocol formats is provided.

The marking of the data packet DP with the priority class KPi assigned to the data packet DP can be carried out in different ways depending on the respective protocol format.

In FIG. 3b, a priority mark PM is entered in the field type for the customary protocol format from FIG. 3a. The priority mark PM is an unambiguous value with which the priority class PKi is unambiguously indicated. This value must generally be reserved unambiguously for the corresponding priority class.

In addition, there is provision in one variant of the method for data packets DP with the lowest priority, that is to say for example of the first priority class PK0, not to be marked at all but rather to be used without modification in accordance with the customary Ethernet Standard.

Since it is possible for an arrangement to operate with layers HL of different levels, for example both with the TCP/IP protocol, on the one hand, and also with the SPX/

IPX protocol, on the other, for different communications connections it is, in one development of the method, advantageous to provide a further field in the data packet DP. The further field is referred to as a flow identifier FID. The flow identifier FID is unambiguously generated in the arrangement transmitting the data packet DP. Thus, it is possible for the receiver of the data packet DP to determine, with reference to the data packet DP, an unambiguous assignment of the data packet DP to the type of the communications connection, the priority mark PM and the Service Access Point, [lacuna] of the coupling to the higher layers. This is possible since, as a result of the procedure described above, the combination of the MAC transmitter address and of the flow identifier FID is unambiguous worldwide, and can thus be used for unambiguous assignment of the data packet DP in the receiver of the data packet DP.

Since, as has been described above, in the case of marked data packets DP the original field content of the field type is described, the information which is contained therein relating to the associated network protocol must previously be transmitted separately to the receiver, for example to the second arrangement A2. This takes place in a signaling phase which is explained in more detail below.

The marking for the protocol format which is described in FIG. 3c is illustrated in FIG. 3d. Here, the priority mark PM is contained in the field DSAP which usually contains the receiver service access point of the higher layer HL.

In addition, the priority mark PM can also be contained, instead or additionally in a field SSAP which usually contains the value of the service access point of the higher layer HL of the transmitter of the data packet DP.

In this variant, the field of the flow identifier FID is, in one development, likewise additionally provided in the data packet DP.

In the format which is illustrated in FIG. 3e, the priority mark PM is, for example, transmitted in the field type (cf. FIG. 3f). In addition, there is, in turn, a field provided for the flow identifier FID. This possibility of converting the priority mark PM, and thus the statement of the priority class PKi, in a marking of the data packet DP constitutes merely one of a large number of examples.

In variants, it is likewise possible to provide new fields for receiving this information.

If different service access points are used for different data streams, the assignment of the data stream to the priority class PKi is written into a table, and each data packet DP which is assigned to the data stream is assigned the corresponding priority class PKi which has been defined in the table.

If, however, it is not the case that a separate access point is provided for each data stream, there is provision to determine for determining the identification number of the port of the data packet DP which has been supplied by a TCP/IP layer. However, other information which identifies the data stream to which the data packet DP is assigned, and which has already been generated in higher layers than, for example, the TCP/IP layer, can also be identified and evaluated.

The result of this is, for example, that in the event that it was possible to identify the associated data stream for the data packet DP, it is possible to perform a mapping onto the corresponding priority class PKi which is to be assigned to the data packet DP.

In one development of the method there is, in addition, provision for the mapping of the quality features which are to be ensured for the data stream, that is to say the communications connection, to be designed so as to be dynamically

configurable during a connection setup phase. If no connection setup phase has been carried out for the datastream to which the data packet DP is assigned, in the simplest case the data packet DP is correspondingly assigned the first priority class PK0.

The arrangement and the method are explained in more detail by means of the following example. The transport protocol used by any desired application ANW (cf. FIG. 4), for example any desired program, uses the UDP/IP protocol in order to transmit video data over the Ethernet.

If the port number of the UDP data stream to be transmitted is known, it is possible for resources to be requested for the connections. Here, for a signaling operation which is described below, a resource administration unit BMV determines which resources are required for which UDP port number.

The resource administration unit BMV carries out, for example, a connection setup phase and marks the new communications connection unambiguously. The marking is detected by the arrangement in that, for example, the UDP port number is evaluated, and, based on this number, the communications connection for the data packet DP and the corresponding priority class PKi can be determined.

FIG. 4 illustrates the arrangement with which the method can, by way of example, be carried out. The arrangement, for example the first arrangement A1, the second arrangement A2 and the further arrangement Ai has at least one data packet classification unit DK, one priority-marking unit PME and a buffer PS.

The arrangement is logically arranged between the higher layers HL and the MAC layer.

The data-packet classification unit DK is configured in such a way that the mapping of the quality features onto priority classes PKi described above can be carried out.

The priority-marking unit PME is configured in such a way that the marking described above, for example therefore the writing of predefined values into specific, predefinable data fields of the data packet DP, is carried out.

The data packet DP to be transmitted is supplied by the higher layers HL of the arrangement. Here, it is fed to the data-packet classification unit DK, which determines a priority class PKi for the data packet DP. The data packet DP is stored in the buffer PS.

In a further development of the method there is provision to assign separate parts of the buffer PS of a predefinable magnitude to the individual priority classes PKi. Only the data packets DP to which the priority classes PKi which corresponds to the part of the buffer PS has been assigned are written in each case into the separate parts which are assigned to the priority classes PKi. The individual parts of the buffer PS are preferably configured according to the First-In-First-Out principle (FIFO buffer). In this way, it is easily possible to implement a so-called queue for the data packets DP before the transmission of the data packets DP.

A simple way of implementing the method is made possible by means of this refinement of the arrangement.

Any desired scheduling method can be used to read out the data packets DP from the parts of the buffer PS. An overview of the various scheduling methods is given in document by H. Zhang and D. Ferrari, Rate-Controlled Static-Priority Queuing, Proc. of INFOCOM '93, San Francisco, Calif., April, 1993.

A very simple algorithm for reading out the data packets DP from the parts of the buffer PS can take the form in which whenever a data packet DP is to be read out of the buffer PS for the transmission of the data packet DP, in each case that part of the buffer PS which corresponds to the highest

priority class, for example to the fourth priority class PK3, is examined and checked to determine whether a data packet DP is waiting for transmission in this part of the buffer PS. If this is not the case, in succession and in descending order of priority, the next part of the buffer PS to which the next lowest priority class PKi-1 is respectively assigned is correspondingly searched for. In one development of the arrangement, a scheduler unit SCH is provided for carrying out the appropriate scheduling method.

In one development of the arrangement, the resource administration unit BMV is additionally provided.

The resource administration unit BMV is configured in such a way that at least one of the following functions is ensured by the resource administration unit BMV:

the resources of the respective arrangement A1, A2 Ai and/or resources of a communications network used for transmitting the data packet DP are administered, these being, for example, the buffer PS, the bandwidth of the Ethernet, which is to be reserved, for example, for a communications connection, the scheduler SCH, etc.; for the data packet DP, an access control is carried out as a function of which it is determined whether or not the data packet DP is further processed;

a signaling operation is carried out between the arrangement and at least one further arrangement with which the arrangement is connected. In this signaling operation, no characteristics of the communications connection, for example the reserved required bandwidth and further quality features for the communications connection which are ensured for the communications connection to which the respective data packet DP is assigned are defined.

In one development of the arrangement it is advantageous to assign a separate, independent Media Access Control address (MAC address) to the resource administration unit BMV. As a result, it becomes possible to implement the resource administration unit independently, only physically connected, of the rest of the arrangement and to communicate with the arrangement only via couplings.

Generation of Artificial Collisions

In one development of the method there is provision, in an arrangement configured as an Ethernet switch, to generate, in the event that the buffer PS or a part of the buffer PS is full and a data packet DP is to be written into the buffer PS or into the filled part of the buffer PS, a so-called artificial collision for the data packet DP to be written in, as a result of which the transmitter is directly informed that the data packet DP cannot be processed at the respective time.

The transmitter of the data packet DP subsequently aborts the transmission of the data packet DP and the transmission attempt of the data packet DP is not repeated until after a randomly selected time period. This prevents data packets DP being lost since the transmitter of the data packet DP completely transmits the data packet DP once more.

However, since in this development of the method, the propagation times for the data packet DP increase by the amount of the time up to the first successful transmission attempt of the data packet DP specifically not only for the data of the communications connections which have to pass through the filled buffer PS but also for all the data packets DP of the braked, i.e. "blocked" transmitter, irrespective of the destination of these data packets DP or of which communications connection these data packets DP belong to, it is advantageous to provide the following development.

However, in the case of data to be transmitted with, for example real-time requirements, excessively long delays are, as has been described above, not acceptable under

11

certain circumstances. For this reason, it is advantageous to protect transmitters with such requirements against artificially generated, undesired collisions of the data packets DP, since it is often more acceptable, with these types of data streams, to accept packet losses which are caused by the overflowing of the buffer PS and/or of the part of the buffer PS than to accept blocking of the entire network interface of the transmitter.

This problem is solved in that the use of the generation of artificial collisions, which is referred to below as backpressure, is permitted, or prohibited, to terminals in a specific and dynamic fashion.

For this purpose, automatic assignment of the statement whether or not the backpressure method is permitted for the respective terminal which is identified by means of the MAC address is implemented. For this reason, it is necessary for the receiver which is configured, for example, as an Ethernet switch, of the data packet DP to be able to detect whether or not the transmitter of the data packet DP is authorized for a backpressure method. This information can be contained, for example, in the packet format which is illustrated in FIGS. 5a to 5f and/or can be entered in a table in the receiver of the data packet DP, in each case in a terminal-specific fashion.

The packet formats (illustrated in FIGS. 3a, 3c and 3e) for an Ethernet data packet DP are illustrated in FIGS. 5a, 5c and 5e.

In this development of the method there is provided, for example in the field type, a status statement BPSTAT with which it is stated whether or not the backpressure method is authorized for the transmitter of the data packet DP.

Likewise, a flow identifier field FID is provided in one development (cf. FIGS. 5b, 5e, 5f). Even if the same fields have been used in this exemplary embodiment, as was described for the priority classes PKi, there is, however, likewise provision in one variant for further fields to be provided in the protocol format, in order thus to be able to transmit both types of information in one data packet DP. Signaling

Possible signaling for the communications connection, in the form in which it is illustrated, for example, in FIG. 1, is described in FIG. 6.

A connection request (Connect Request) is transmitted to the second arrangement A2 from the first arrangement A1. The Connect Request contains, for example, the target address DA, the transmitter address SA, the flow identifier statement FID, the field type as well as a parameter field TSpec in which it is stated which quality features are requested for the communications connection, as well as, for example, a service class field RSpec, in which it is indicated which service class, for example, which priority class PKi is requested for the communications connection.

The Connect Request is passed on directly to the further arrangement Ai from the second arrangement A2. In the case of a communications connection which is routed via a plurality of Ethernet switching units, the second arrangement A2 is to be understood as a quantity of switching units which each pass on the Connect Request and the corresponding further data packets DP to that terminal which constitutes the corresponding terminal for the communications connection.

The terminal, for example the further arrangement Ai, transmits a Connect Reply which, in addition to the target address field DA, the transmitter address SA, the flow identifier FID, contains a Result (Success) statement, in which it is stated whether or not the requested communications connection has been accepted.

In addition, in a Reason result statement it is stated why the communications connection has, for example, not been

12

accepted. A cause of this can, for example, be an excessively low available bandwidth in the communications network. These statements can be coded in a predefinable way, and these fields thus contain only random numerical values.

The Connect Reply response is passed onto the first arrangement A1 via "all the second arrangements" A2. After the reception of the Connect Reply response in the first arrangement A1 and of a positive response, i.e. a response to the effect that the communications connection has been accepted, this connection setup phase described above is terminated.

This is followed by the actual transmission of the useful data, that is to say of the data packets DP to which in each case there is assigned according, for example, to the negotiated priority class PKi or even a freely predefinable priority class PKi which can be defined by the transmitter [lacuna], or else which are in each case assigned adaptively to the data packet DP, for example as a function of the network loading of the respective Ethernet switching arrangement.

If a communications connection is aborted, this is achieved, for example, by virtue of the fact that a connection abort request (Disconnect Request) is transmitted to the second arrangement A2 by a transmitter, for example the first arrangement A1, and is passed on from the second arrangement A2 to the further arrangement Ai. The connection abort request Disconnect Request contains, for example, the target address DA, the transmitter address SA and the flow identifier FID.

Implementing the Priority-controlled Transmission at the Bit Transmission Layer

In a development of the method there is provision for the bandwidth which is available, in each case, for example in the Ethernet or else for a communications connection within the scope of the Ethernet, is divided into a plurality of bandwidth areas, for example corresponding to the number of priority classes PKi provided.

Each part of the bandwidth serves to transmit the data packets DP which are assigned to the priority class PKi to which the respective part of the bandwidth is also assigned.

The quite simple case in which only two priority classes PK0 and PK1 are provided (cf. FIG. 7) is represented below.

For this case, for example, two time slots, a first time slot ZS1 and a second time slot ZS2, are provided. The first time slot ZS1 is, for example, for high-priority data, that is to say for data packets DP to which the second priority class PK1 has been assigned, and the second time slot ZS2 is provided for data packets to which the low priority class PK0 has been assigned.

Here, there is provision, in a development of the method, to simplify the method in such a way that only the magnitude of the first time slot ZS1 is administered. This means that it is advantageous for the time slots ZS1, ZS2 to be of variable configuration in terms of their width. Starting out from a freely predefinable reference time t_0 , a starting time of the second time slot ZS2 is obtained from the sum of the reference time t_0 and a time slot length S1 of the first time slot ZS1.

The time slot length S1 is selected in such a way that offered traffic A_H of all the high-priority data streams can be transmitted. The following applies:

$$A_H < \frac{S1}{S_F} \cdot B$$

B designates the bandwidth of the corresponding Ethernet segment. S_F designates the sum of the time slot length S1 of the first time slot and a time slot length S2 of the second time

13

slot ZS2 ($S_F = S1 + S2$). The above rule gives only a rough guide to the dimensioning of the first time slot length S1. Any desired rule for dimensioning, i.e. dividing up of the individual time slot lengths is known to the person skilled in the art.

If there are requests for the maximum delays T_{Dmax} of the data which are to be transmitted with high priority, the following rule, for example, is to be complied with:

$$S_F < \min T_{Dmax} \text{ (all the high-priority communications + S1 connections)}$$

Furthermore, in one development of the method it is advantageous to synchronize the reference time to of a freely predefinable arrangement, for example by means of timing messages at freely predefinable time intervals, i.e. the reference time t_0 is transmitted to all the other arrangements by an arrangement in freely predefinable time intervals. The timing message is received by the arrangements and used as their new reference time t_0 . Genuine synchronization of the arrangements is not possible in the Ethernet. However, this periodic freshening up of the reference times t_0 permits sufficiently precise coordination of the clocks contained in the operating systems of the arrangements.

The invention is not limited to the particular details of the method and apparatus depicted and other modifications and applications are contemplated. Certain other changes may be made in the above described method and apparatus without departing from the true spirit and scope of the invention herein involved. It is intended, therefore, that the subject matter in the above depiction shall be interpreted as illustrative and not in a limiting sense.

What is claimed is:

1. A method of transmitting a data packet in the Ethernet from a first arrangement to at least a second arrangement, comprising the steps of:

assigning a data packet, that is to be transmitted, is one of at least two priority classes;

unambiguously marking the data packet according to the priority class, whereby for a respective data packet an access control is carried out taking into account the priority class assigned to the respective data packet, the respective data packet being transmitted to the second arrangement in accordance with a predefinable sequence which is specified by a predefinable scheduling method; and

wherein, in the event that the buffer or at least part of the buffer is filled with stored data packets which have not yet been read out again, and a further data packet to be transmitted is to be stored in the buffer or at least in part of the buffer, an artificial collision is produced and communicated to the arrangement transmitting the data packet.

2. A method of transmitting a data packet in the Ethernet from a first arrangement to at least a second arrangement, comprising the steps of:

assigning a data packet, that is to be transmitted, is one of at least two priority classes;

unambiguously marking the data packet according to the priority class, whereby for a respective data packet an access control is carried out taking into account the priority class assigned to the respective data packet, the respective data packet being transmitted to the second arrangement in accordance with a predefinable sequence which is specified by a predefinable scheduling method; and

wherein an artificial collision is produced only if the data packet contains an item of information specifying that

14

the data packet is permitted to produce an artificial collision for the arrangement transmitting the data packet.

3. A method of transmitting a data packet in the Ethernet from a first arrangement to at least a second arrangement, comprising the steps of:

assigning a data packet, that is to be transmitted, is one of at least two priority classes;

unambiguously marking the data packet according to the priority class, whereby for a respective data packet an access control is carried out taking into account the priority class assigned to the respective data packet, the respective data packet being transmitted to the second arrangement in accordance with a predefinable sequence which is specified by a predefinable scheduling method; and

wherein at least one time message is generated by a predetermined arrangement at predefinable time intervals and is transmitted to at least some of further arrangements which are coupled to the arrangement,

wherein the time message contains a synchronization message, as a function of which time intervals are determined which are provided for transmitting data packets of a respective predefinable priority class,

wherein the time message is received from at least some of the further arrangements,

wherein the received time message is evaluated, the synchronization time being determined, as a function of which time intervals are determined which are provided for transmitting data packets of a respective predefinable priority class.

4. The method as claimed in claim 3, wherein the data packet is transmitted only in the time interval which is provided for the priority class which has been assigned to the data packet.

5. An arrangement for transmitting a data packet in the Ethernet from a first arrangement to at least a second arrangement: comprising:

a data-packet classification unit with which one priority class of at least two priority classes is assigned to a data packet to be transmitted;

a priority marking unit with which the data packet is unambiguously marked in accordance with the priority class;

a buffer in which at least the data packet to be transmitted is buffered; and

a resource administration unit which configured such that an access control for each data packet is carried out for taking into account the priority class assigned to the respective data packet, the respective data packet being transmitted to the second arrangement according to a predefinable sequence which is specified by the predefinable scheduling method.

6. The arrangement as claimed in claim 5, wherein the arrangement further comprises:

a receiver unit with which a data packet to be transmitted is received in an Ethernet switching unit, and

a priority analysis unit with which a priority class assigned to the data packet is determined.

7. The arrangement as claimed in claim 5 wherein, the resource administration unit is structured such that at least one of the following functions is ensured:

administration of resources of at least one of the arrangement and resources of a data network used for transmitting the data packet,

15

a signaling operation between the arrangement and at least a further arrangement to which the arrangement is coupled with which characteristics of a communications link to which the data packet is assigned are specified.

8. The arrangement as claimed in claim 7, wherein the resource administration unit has a Media Access Control address.

9. The arrangement as claimed in claim 5, wherein at least some of the buffer is assigned to a priority class and only a data packet which has the corresponding priority class is stored.

16

10. The arrangement as claimed in claim 5, wherein at least one of the buffer and at least a part of the buffer is a First-In-First-Out buffer.

5

* * * * *

10